arXiv:1704.04054v1 [cs.CV] 13 Apr 2017

Saliency-guided Adaptive Seeding for Supervoxel Segmentation

Ge Gao, Mikko Lauri and Simone Frintrop

***Abstract*— We propose a new saliency-guided method for generating supervoxels in 3D space. Rather than using an evenly distributed spatial seeding procedure, our method uses visual saliency to guide the process of supervoxel generation. This results in densely distributed, small, and precise supervoxels in salient regions which often contain objects, and larger supervoxels in less salient regions that often correspond to background. Our approach largely improves the quality of the resulting supervoxel segmentation in terms of boundary recall and under-segmentation error on publicly available benchmarks.**

I. INTRODUCTION

Superpixels are currently a popular way of grouping pixels to larger entities and thus to strongly reduce the number of primitives that subsequent modules have to deal with [21]. This is especially important in robotics scenarios, where real-time constraints usually make the full interpretation of high resolution image data infeasible. A popular application of superpixels is object discovery, where superpixels are grouped to obtain object candidates which in turn are used as input for object recognition methods [14], [9], [10]. Other applications include automatic object handle grasping [20], unknown object manipulation in cluttered environments [3], and semantic segmentation [24].

Since superpixels serve as input to all further processing, their quality has a significant impact on the quality of the output [8]. For example, violating object boundaries will introduce a permanent error into the processing pipeline since all the following algorithms will be forced to use superpixels which contain more than one object [16]. In other words, if a single superpixel contains pixels belonging to two distinct objects, the two objects can not be fully distinguished in any subsequent processing step. The quality of a set of superpixels can be assessed based on, e.g., their adherence to object boundaries, compactness, smoothness, and controllable number of superpixels [21]. Among these qualities, boundary adherence is one of the most important requirements due to the above mentioned reasons.

While dozens of superpixel algorithms exist that operate on 2D images, see, e.g., [1], [5], [22], [4], [27] for some examples, methods that analyze RGB-D data and generate so called supervoxels are much less common. Among the few methods that exploit depth, the Voxel Cloud Connectivity Segmentation (VCCS) supervoxels presented by Papon et al. [16] is probably the best known and most used method[1]. VCCS oversegments a 3D point cloud into supervoxels by applying a spatial seeding procedure that places the initial supervoxel centroids on a regular grid. The main difficulty of the method, as well as of other superpixel and supervoxel methods, is to find an appropriate parameter setting that defines a reasonable tradeoff between the number of supervoxels and the precision. On one hand, obtaining a lower number of larger superpixels is favorable since this reduces the computational complexity of subsequent processing. On the other hand, the superpixels should not be larger than the smallest objects, since otherwise, these will not be properly segmented and will be permanently lost in the further process [1].

Fig. 1. An example scene for supervoxel computation. Top left: original image; top right: the corresponding saliency map. Middle left: 2D projection of VCCS supervoxels which are uniformly distributed. Middle right: 2D projection of our saliency-guided supervoxels (SSV) with adaptive resolution. Bottom left: 3D supervoxel cloud for VCCS. Bottom right: 3D supervoxel cloud for our SSV method. Our method has a sparser supervoxel structure in non-salient background regions and a denser structure in highly salient regions that often contain objects.

In this paper, we present an improved supervoxel algorithm called Saliency-guided Supervoxels (SSV) which is based on VCCS [16]. Instead of seeding the initial supervoxel centroids evenly spaced over the whole input data, we propose a saliency-based regional seeding procedure as illustrated in Fig. 1. This means, we adaptively seed the supervoxels guided by saliency: highly salient regions will have a greater density of seeds, and vice versa. To achieve this, we cluster the input data according to visual

Department of Informatics, University of Hamburg, 22527 Hamburg, Germany {gao, lauri, frintrop}@informatik.uni-hamburg.de

[1]The code for VCCS is publicly available as part of the Point Cloud Library (PCL).

saliency, assigning to each cluster a saliency-specific seeding resolution. This leads to small, precise supervoxels in salient regions, and to large supervoxels in less salient ones, for example on homogeneous background surfaces (see e.g. the wall in Fig. 1, bottom right).

The idea is motivated by the fact that saliency highlights regions that visually stick out of the image, i.e., they differ according to some features — such as color, intensity, or depth — from their background [6]. Thus, saliency is also a good indicator for the presence of objects, since objects also usually differ from their surroundings [14]. Here, this property is especially useful since it enables us to obtain a higher density of supervoxels in regions that probably correspond to objects. We show that, with the same average number of supervoxels per image as VCCS, we get a clearly improved boundary recall, undersegmentation error, and explained variation of the supervoxel segmentation. We thus show that visual saliency provides a useful prior that allows the segmentation to better respect object boundaries.

The organization of the paper is as follows. In Section II, we review related work. In Section III, we give an overview of our proposed method, as well as the details of each step. We experimentally evaluate our method in Section IV on two publicly available datasets. Section V provides some concluding remarks.

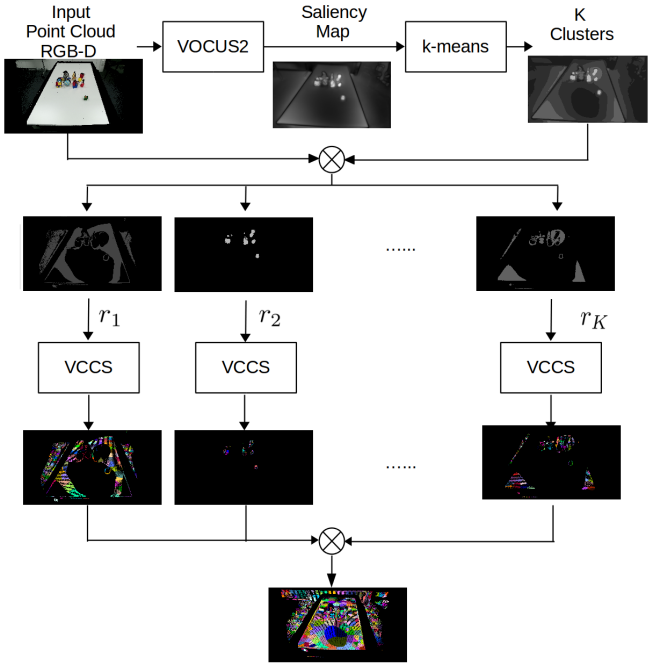

Fig. 2. System overview. First, the saliency system VOCUS2 [6] generates a saliency map. The point cloud is partitioned into K clusters using k-means on the saliency values, and each cluster k is assigned a seeding resolution r_k. The VCCS supervoxel segmentation method [16] is applied to each cluster independently and the results are combined to form the final output.

II. BACKGROUND AND RELATED WORK

Following the classification of Stutz et al. [21], superpixel methods can be classified into watershed-based [12], density-based [22], graph-based [5], contour evolution [11], path-based [7], clustering-based [1], [23] and energy optimization methods [4], [27]. Among the most popular methods is the Simple Linear Iterative Clustering (SLIC) superpixel method [1] that applies local k-means clusterings of the image pixels on a regular grid pattern over the entire image to generate perceptually uniform superpixels. SLIC is appealing since it is simple, fast, and has only two parameters that must be controlled: the number of superpixels that shall be generated and the compactness of the superpixels. Many other superpixel and supervoxel methods follow the idea of SLIC, e.g., [23], [26], and [16].

Few methods integrate depth data into the segmentation process. The Depth-Adaptive Superpixels (DASP) [23] extend the iterative local clustering approach by introducing the control of superpixel seed density based on depth information. Similarly, Yang et al. [26] enhance the superpixel generation process using the depth difference between pixels to prevent violating boundaries between objects of similar color. They also use a local k-means clustering approach with an eight-dimensional distance measure including color, 2D and 3D coordinates. Both methods can be classified as 2.5D methods, since they use depth to improve the superpixel generation, but they do not create 3D supervoxels but rather 2D superpixels.

Instead, the Voxel Cloud Connectivity Segmentation (VCCS) [16] generates supervoxels from RGB-D data, while guaranteeing that all voxels within each supervoxel remain spatially connected. Similarly to SLIC and DASP, VCCS is also a variant of iterative local clustering applied on a regular lattice. But in contrast to these methods, VCCS takes advantage of the 3D geometry of the scene and generates a full 3D supervoxel graph.

Our Saliency-guided Supervoxel (SSV) method is based on VCCS, but instead of uniformly distributing the supervoxels over the data, we use an adaptive seeding procedure guided by saliency. Our results show that our method clearly improves the quality of the supervoxels according to several evaluation metrics.

It should be mentioned that in the literature other methods are also referred to as supervoxels, which use time instead of space as the third dimension and are thus a type of video segmentation [25], [7]. These methods are not applicable to RGB-D point clouds since they usually require a solid 3D volume of time and space as input. Thus, such methods are not comparable with the supervoxel methods operating on 3D data from a sensor such as the Microsoft Kinect.

III. SALIENCY-GUIDED SUPERVOXEL SEGMENTATION

In this section, we present Saliency-guided Supervoxels (SSV), our algorithm for generating supervoxels using saliency-guided adaptive seeding. SSV combines a visual saliency model, a newly introduced saliency-guided adaptive seeding approach and a supervoxel segmentation algorithm.

An overview of the algorithm is presented in Section III-A. We then give more details on the visual saliency model and supervoxel computation in Sections III-B and III-C,

respectively. The adaptive seeding approach is described in Section III-D.

A. System overview

As illustrated in Figure 2, the input to SSV is an RGB-D point cloud, as obtained from an RGB-D camera such as the Microsoft Kinect. During pre-processing, the RGB information is passed to a visual saliency system, here VOCUS2 [6], to generate a pixel-level saliency map. The pixel level saliency map is segmented into K regions applying k-means. The same segmentation is applied to partition the point cloud into K clusters. The clusters are sorted in an ascending order based on their average saliency values.

The supervoxel seed resolution is determined according to the average saliency value: clusters with a higher average saliency have denser seeding. As supervoxel method we use VCCS [16], which is applied to the points in each cluster independently. Finally, the results are merged to obtain the final supervoxel segmentation. Since setting $K = 1$ reduces our method to be equivalent to VCCS, SSV may be viewed as a generalization of VCCS.

B. Saliency model

To compute a pixel-level saliency map from the RGB data, we use the computational visual saliency method VOCUS2 [6]. VOCUS2 converts its input image into an opponent-color space with intensity, red-green, and blue-yellow color channels. Center-surround contrasts are then computed on multiple scales by Difference-of-Gaussian filters operating on center and surround twin pyramids. Finally, the saliency map is generated by fusing the contrast maps across scales and channels using a simple arithmetic mean.

We use VOCUS2 since it is fast and has obtained good results on several benchmarks. Furthermore, it does not have a center-bias, which is important in robotic applications since objects of interest are usually not in the center of the image. Despite these considerations, any computational visual saliency model could be applied in SSV.

C. Voxel cloud connectivity segmentation

VCCS is a supervoxel segmentation algorithm introduced by Papon et al. [16]. VCCS first voxelizes the input RGB-D point cloud by equally partitioning the 3D space using an octree structure. The size of each voxel is defined by the voxel resolution R_{voxel}.

Supervoxel seeds are generated uniformly on a regular grid with resolution R_{seed} that is much larger than R_{voxel}. Higher R_{seed} results in fewer supervoxels and vice versa. For each occupied seed voxel, the nearest cloud voxel is selected as an initial seed. Unoccupied seed voxels are discarded.

After selecting the initial seeds, a local clustering of voxels is performed iteratively until all the voxels are assigned to supervoxels. The clustering is performed in a 39 dimensional feature space:

$$\mathbf{F} = [x, y, z, L, a, b, \text{FPFH}_{1\ldots33}], \tag{1}$$

where x, y, z are spatial coordinates, L, a, b denote color in CIELab space, and $\text{FPFH}_{1\ldots33}$ are the 33 elements of Fast Point Feature Histogram [17]. Each voxel is assigned to the supervoxel to whose centroid it has the smallest normalized distance D. The distance measure D is

$$D = \sqrt{\frac{\lambda {D_c}^2}{m^2} + \frac{\mu {D_s}^2}{3{R_{seed}}^2} + \epsilon {D_{\text{HiK}}}^2}, \tag{2}$$

where D_c is the Euclidean distance in CIELab color space, D_s denotes the spatial distance, D_{HiK} is the Histogram Intersection Kernel of FPFH [2], m is a normalization constant, and λ, μ and ϵ are the weights for each distance such that $\lambda + \mu + \epsilon = 1$.

D. Saliency-guided adaptive seeding

For our saliency-guided supervoxel seeding, we first compute a saliency map as described above. We then use k-means to partition the pixel-level saliency map into K clusters.

Let C_k denote the set of pixels in the k-th cluster, and denote by x_i the saliency value of a pixel $i \in C_k$. The average saliency of cluster k is given by

$$sal_k = \frac{1}{|C_k|} \sum_{i \in C_k} x_i. \tag{3}$$

We sort the clusters in ascending order according to the average saliency.

To control the size of supervoxels, the minimum and maximum seed resolution R_{min} and R_{max}, respectively, are defined. A step size d is determined by

$$d = -\frac{\log R_{min} - \log R_{max}}{K - 1}. \tag{4}$$

The seed resolution r_k for the k-th cluster ranked according to average saliency is

$$r_k = 10^{\log R_{max} - (k-1)d}, \tag{5}$$

i.e., regions with a high average saliency are seeded densely while less salient regions have a sparse seed distribution.

We apply VCCS (Section III-C) independently to each cluster, using r_k as the seed resolution R_{seed} for data in cluster k. The final result is obtained by combining the K supervoxelizations into a single multiple seed resolution supervoxel representation. An example of the resulting 3D supervoxelization is shown in Figure 3, together with the VCCS supervoxelization for comparison.

IV. EXPERIMENTAL RESULTS AND EVALUATION

We compare our supervoxel method, SSV, against the current state-of-the-art supervoxel method VCCS [16]. The evaluation details and metrics we apply are described in Sec. IV-A, the datasets in Sec. IV-B, and the experimental results in Sec. IV-C.

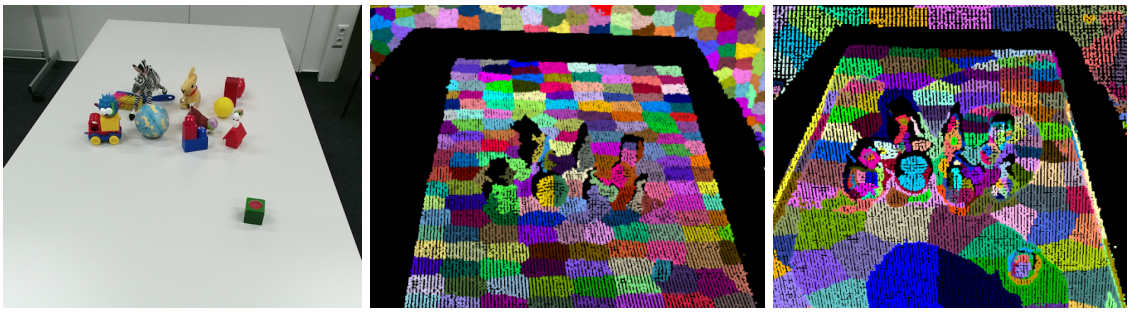

Fig. 3. Input image (left) and supervoxels obtained with VCCS (middle) and with our SSV method (right).

A. Evaluation framework

We compare SSV and VCCS using the superpixel benchmarking tool[2] of Stutz et al. [21]. The tool provides a platform to evaluate the performance of superpixel algorithms using common metrics such as boundary recall (REC), undersegmentation error (UE), and explained variation (EV). The benchmark tool evaluates supervoxel methods by projecting each supervoxel back to the 2D image plane, and then evaluating the result similarly as superpixel methods. This is reasonable since most publicly available datasets provide the ground truth only as a 2D image.

The notation we use in this section is as follows. Given an input image $I = \{x_j\}_{j=1}^{J}$ with J pixels, we write $S = \{S_n\}_{n=1}^{N}$ to denote a segmentation of I into N superpixels, and $G = \{G_m\}_{m=1}^{M}$ to denote the M ground truth segments.

Boundary recall (REC) [13] assesses boundary adherence with ground truth by measuring the percentage of the superpixel edges that fall within a certain range of an arbitrary ground truth boundary. The range is defined as $(2r+1)^2$, with $r = 0.0025 \times L$ where L is the image diagonal size [21]. Higher boundary recall indicates a better boundary adherence.

Undersegmentation error (UE) [11] measures the amount of "leakage" of a segmentation S with respect to the ground truth G:

$$\mathrm{UE}(S,G) = \frac{1}{M}\sum_{m=1}^{M}\frac{\left[\sum_{\{n|S_n\cap G_m\neq\emptyset\}}|S_n|\right]-|G_m|}{|G_m|}, \quad (6)$$

where the leakage of superpixel S_n with respect to ground truth G_m is represented by the inner term. Lower undersegmentation error demonstrates less leakage.

Explained variation (EV) [15] attempts to measure the quality of a superpixel segmentation without relying on a human-annotated ground truth. EV is defined as

$$\mathrm{EV}(S) = \frac{\sum_{n=1}^{N}|S_n|\left(\mu(S_n)-\mu(I)\right)^2}{\sum_{j=1}^{J}\left(x_j-\mu(I)\right)^2}, \quad (7)$$

where x_j denotes the actual value for pixel j, $\mu(I)$ is the global pixel mean and $\mu(S_n)$ is the mean value of superpixel n. EV denotes the proportion of image variation that can be explained by superpixel segments. A higher EV indicates better performance.

[2]https://github.com/davidstutz/superpixel-benchmark

B. Datasets

We evaluate on two publicly available RGB-D datasets, NYUV2 [18] and SUNRGBD [19]. NYUV2 contains 1449 color images with associated depth information. The data are collected applying Microsoft Kinect v1, and depict varying indoor scenes. We use the test set of 399 randomly chosen images specified in the superpixel benchmark to evaluate our method.

SUNRGBD contains 10335 images of cluttered indoor scenes, likewise with color and depth information. The dataset contains images collected applying the Intel RealSense, Asus Xtion, Microsoft Kinect v1 and v2 sensors. We use 2000 randomly chosen images from this dataset to evaluate our method.

C. Experimental results

During our experiments, we varied the seed resolution R_{seed} between 0.05 and 0.25 for VCCS. For SSV, we created between 2 and 6 clusters, with minimum seed resolution R_{min} between 0.05 and 0.19 and maximum seed resolution $R_{max} = 0.4$. The values were chosen in order to obtain approximately the same average number of superpixels for both methods[3]. For both SSV and VCCS, we set $R_{voxel} = 0.008$, color weight $\lambda = 0.7$, spatial distance weight $\mu = 0.15$ and HiK distance weight $\epsilon = 0.15$ (Equation (2)).

Figure 4 shows the REC, UE and EV for SSV and VCCS as a function of the average number of superpixels. For this figure, we show data corresponding to seed resolution R_{seed} of VCCS is between 0.09 and 0.18. For SSV, we set $K = 6$, $R_{max} = 0.2$, and R_{min} between 0.06 and 0.19. The means and their 95% confidence intervals are plotted. The figure shows clearly that with the same amount of superpixels, SSV achieves significantly better performance in terms of REC, UE, and EV. This is especially true for smaller numbers of superpixels (left part of plots), which is of special interest for many applications that require low complexity. Due to the saliency prior, supervoxels generated by SSV are more

[3]For both methods, it is not possible to determine the number of superpixels precisely in advance, due to the projection from 3D.

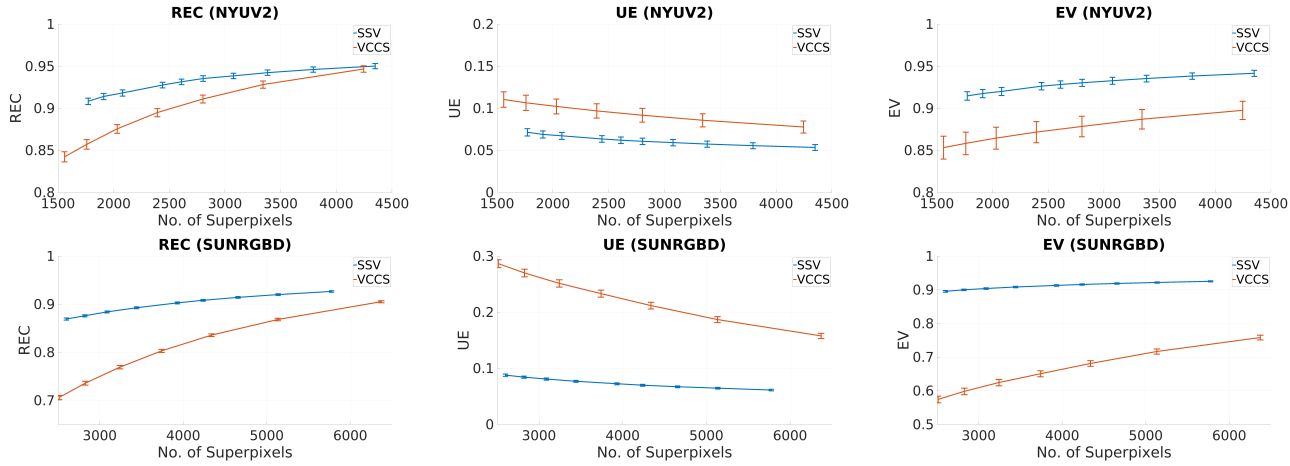

Fig. 4. Average boundary recall (REC, higher is better), undersegmentation error (UE, lower is better) and explained variation (EV, higher is better) with 95% confidence intervals in the NYUV2 (top row) and SUNRGBD (bottom row) datasets. SSV clearly outperforms VCCS with a similar number of superpixels.

efficient: a larger fraction of the supervoxels are in regions with many boundaries to preserve, while fewer supervoxel are used in background regions.

Table I indicates how the performance of VCCS and SSV varies as a function of the key parameters: R_{seed} for VCCS, and K, R_{min}, and R_{max} for SSV. We note a general and expected trend that a higher number of superpixels leads to higher boundary recall and explained variation, as well as lower undersegmentation error. However, more superpixels also corresponds to less reduction in the complexity of the input representation.

In Figure 5, we show a qualitative comparison of our method and VCCS. The fact that SSV uses saliency to guide the seeding process is seen from having a higher density of supervoxels in highly salient areas that are likely to contain objects. Supervoxels are generated with a low density in non-salient background areas such as walls, resulting in larger supervoxels. We also see that small objects are captured much better with SSV, note for example the details of the sink (handles and tap) on the 3rd row.

For both datasets, the average processing time of VCCS is 1.3 sec. per image. For SSV, the average processing time increases with the number of k-means clusters. With $K = 5$ clusters, the average processing time of SSV is 3.5 sec. per image, and with $K = 6$ clusters, 4 sec. per image with non-optimized code. We remark that SSV is easy to parallelize by computing the supervoxels for each cluster in a separate thread, which we expect will considerably speed up the approach.

V. CONCLUSION

Ensuring the high quality of an initial segmentation is important in applications that require precise object boundaries, e.g., manipulation, since any errors in segmentation propagate through the whole image processing pipeline, and typically cannot be resolved in later processing stages.

Motivated by this observation, we propose to apply visual saliency to guide the oversegmentation of an RGB-D point cloud. More supervoxels are generated in highly salient regions that are likely to contain objects, while less supervoxels are generated in low-saliency background regions.

As our experiments show, our approach preserves object boundaries significantly better than the current state-of-the-art supervoxel method which generates supervoxels uniformly over the whole scene. In our future work, we will investigate the applicability of our supervoxel segmentation method in vision tasks such as object discovery. We are also preparing to release an implementation of our segmentation method publicly as open-source software.

REFERENCES

[1] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Süsstrunk. SLIC superpixels compared to state-of-the-art superpixel methods. *IEEE TPAMI*, 34(11):2274–2282, 2012.

[2] A. Barla, F. Odone, and A. Verri. Histogram intersection kernel for image classification. In *ICIP*, 2003.

[3] A. Boularias, J. Bagnell, and A. Stentz. Learning to manipulate unknown objects in clutter by reinforcement. In *AAAI*, 2015.

[4] M. Van den Bergh, X. Boix, G. Roig, B. de Capitani, and L. Van Gool. SEEDS: superpixels extracted via energy-driven sampling. In *ECCV*, 2012.

[5] F. Felzenszwalb and D. Huttenlocher. Efficient graph-based image segmentation. *IJCV*, 59(2):167–181, 2004.

[6] S. Frintrop, T. Werner, and G. Martín García. Traditional saliency reloaded: A good old model in new shape. In *CVPR*, 2015.

[7] H. Fu, X. Cao, D. Tang, Y. Han, and D. Xu. Regularity preserved superpixels and supervoxels. *IEEE Trans. on Multimedia*, 16(4):1165–1175, 2014.

[8] A. Hanbury. How do superpixels affect image segmentation? In *CIARP*, 2008.

[9] J. Hosang, R. Benenson, P. Dollár, and B. Schiele. What makes for effective detection proposals? *IEEE TPAMI*, 38(4):814–830, 2015.

[10] A. Kanezaki and T. Harada. 3D selective search for obtaining object candidates. In *IROS*, 2015.

[11] A. Levinshtein, A. Stere, K. Kutulakos, D. Fleet, S. Dickinson, and K. Siddiqi. Turbopixels: Fast superpixels using geometric flows. *IEEE TPAMI*, 31(12):2290–2297, 2009.

[12] V. Machairas, M. Faessel, D. Cárdenas-Peña, T. Chabardes, T. Walter, and E. Decencière. Waterpixels. *IEEE TIP*, 24(11):3707–3716, 2015.

TABLE I

EFFECT OF PARAMETER CHOICES ON VCCS AND SSV PERFORMANCE ON THE NYUV2 DATASET. MEAN ± 95% CONFIDENCE INTERVAL FOR THE NUMBER OF SUPERPIXELS (#SP), BOUNDARY RECALL (REC), UNDERSEGMENTATION ERROR (UE), AND EXPLAINED VARIATION (EV).

	R_{seed}		#SP	REC	UE	EV
VCCS	0.01		54806±2889	0.999±0.0004	0.01±0.0005	0.98±0.002
	0.05		13043±770	0.99±0.002	0.04±0.003	0.95±0.006
	0.1		3341±183	0.93±0.004	0.07±0.008	0.89±0.01
	0.15		1558±83	0.84±0.006	0.11±0.009	0.85±0.01
	0.2		914±45	0.77±0.008	0.13±0.01	0.83±0.01
	K	R_{min}	R_{max}	#SP	REC	UE	EV
SSV	6	0.1	0.20	2803±123	0.94±0.003	0.06±0.004	0.93±0.004
			0.25	2370±104	0.92±0.004	0.06±0.004	0.92±0.005
			0.30	2071±91	0.91±0.004	0.07±0.004	0.92±0.005
			0.35	1857±81	0.90±0.004	0.08±0.005	0.91±0.005
			0.40	1704±72	0.89±0.005	0.08±0.005	0.91±0.005
	6	0.01	0.2	17272±1064	0.97±0.002	0.004±0.002	0.97±0.002
		0.05		5170±265	0.95±0.003	0.05±0.003	0.95±0.003
		0.10		2803±123	0.94±0.003	0.06±0.004	0.93±0.004
		0.15		2080±85	0.92±0.004	0.07±0.004	0.92±0.005
		0.19		1771±70	0.91±0.004	0.07±0.004	0.91±0.005
	2	0.1	0.2	2138±113	0.86±0.006	0.09±0.007	0.9±0.008
	3			2243±111	0.87±0.005	0.08±0.006	0.91±0.006
	4			2428±115	0.91±0.004	0.07±0.005	0.92±0.006
	5			2628±120	0.92±0.004	0.07±0.004	0.93±0.005
	6			2803±123	0.94±0.003	0.06±0.004	0.93±0.004

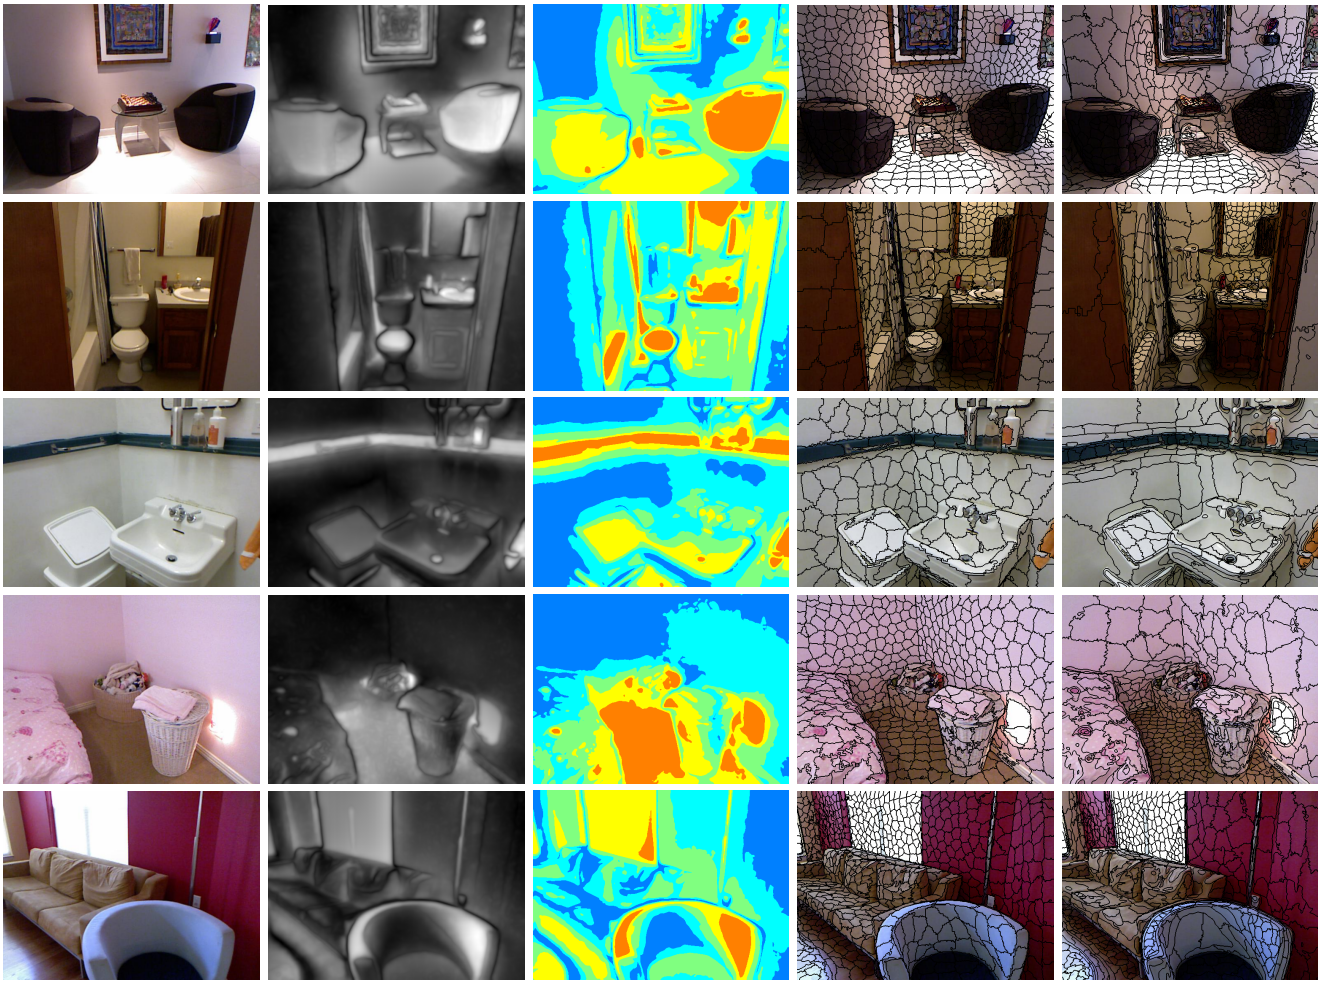

Fig. 5. A qualitative comparison of SSV and VCCS. From left to right: input image, saliency map, result of k-means clustering, supervoxels from VCCS, and our SSV supervoxels.

[13] D. Martin, C. Fowlkes, and J. Malik. Learning to detect natural image boundaries using local brightness, color, and texture cues. *IEEE TPAMI*, 26(5):530–549, 2004.
[14] G. Martín García, E. Potapova, T. Werner, M. Zillich, M. Vincze, and S. Frintrop. Saliency-based object discovery on RGB-D data with a late-fusion approach. In *ICRA*, 2015.
[15] A. Moore, S. Prince, J. Warrell, U. Mohammed, and G. Jones. Superpixel lattices. In *CVPR*, 2008.
[16] J. Papon, A. Abramov, M. Schoeler, and F. Wörgötter. Voxel cloud connectivity segmentation - supervoxels for point clouds. In *CVPR*, 2013.
[17] R. Rusu, N. Blodow, and M. Beetz. Fast point feature histograms (FPFH) for 3D registration. In *ICRA*, 2009.
[18] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus. Indoor segmentation and support inference from RGBD images. In *ECCV*, 2012.
[19] S. Song, S. P. Lichtenberg, and J. Xiao. SUN RGB-D: A RGB-D scene understanding benchmark suite. In *CVPR*, June 2015.
[20] S. Stein, F. Wörgötter, M. Schoeler, J. Papon, and T. Kulvicius. Convexity based object partitioning for robot applications. In *ICRA*, 2014.
[21] D. Stutz, A. Hermans, and B. Leibe. Superpixels: An evaluation of the state-of-the-art. *arXiv preprint arXiv:1612.01601*, 2016.
[22] A. Vedaldi and S. Soatto. Quick shift and kernel methods for mode seeking. In *ECCV*, 2008.
[23] D. Weikersdorfer, D. Gossow, and M. Beetz. Depth-adaptive superpixels. In *ICPR*, 2012.
[24] D. Wolf, J. Prankl, and M. Vincze. Enhancing semantic segmentation for robotics: the power of 3-D entangled forests. *IEEE RAL*, 1(1):49–56, 2016.
[25] C. Xu and J. J. Corso. Evaluation of super-voxel methods for early video processing. In *CVPR*, 2012.
[26] J. Yang, Z. Gan, K. Li, and C. Hou. Graph-based segmentation for RGB-D data using 3-D geometry enhanced superpixels. *IEEE Trans. on Cybernetics*, 45(5):927–940, 2015.
[27] J. Yao, M. Boben, S. Fidler, and R. Urtasun. Real-time coarse-to-fine topologically preserving segmentation. In *CVPR*, 2015.